

Maximum a Posteriori on a Submanifold: a General Image Restoration Method with GAN

Fangzhou Luo
Dept. of Elec & Comp Engineering
McMaster University
Hamilton, Canada
fluo1993@gmail.com

Xiaolin Wu
Dept. of Elec & Comp Engineering
McMaster University
Hamilton, Canada
xwu@ece.mcmaster.ca

Abstract—We propose a general method for various image restoration problems, such as denoising, deblurring, super-resolution and inpainting. The problem is formulated as a constrained optimization problem. Its objective is to maximize a posteriori probability of latent variables, and its constraint is that the image generated by these latent variables must be the same as the degraded image. We use a Generative Adversarial Network (GAN) as our density estimation model. Convincing results are obtained on MNIST dataset.

Index Terms—image restoration, deep learning, optimization

I. INTRODUCTION

Image restoration has been researched for many years, but in a case-by-case way [1]–[5]. Almost all image restoration algorithms are only designed for a certain type of images or degradation. This research paradigm has some obvious disadvantages. It is exhausting to invent new algorithms or train new models for slightly different situations. Even if we can, those specialized solutions are not so elegant, because they are very unlike one another even though the problems they focus on are fundamentally so similar.

It is worth noting that any image degradation process can be abstracted as a many-to-one function. More specifically, for any given degradation process, one degraded image could be degraded from many possible original images. From that point of view, we propose a general method for various image restoration problems, such as denoising, deblurring, super-resolution and inpainting. Our algorithm chooses the most probable original image from all those possible original images, and uses it as the restoration of the given degraded image. To be more precise, the general image restoration is formulated as a constrained optimization problem. Its objective is to maximize a posteriori probability of latent variables, and its constraint is that the image generated by these latent variables must be the same as the degraded image.

Recent progress of density estimation techniques has made our algorithm possible. In the field of image generation, Generative Adversarial Networks (GANs) has achieved great success in recent years [6]–[8]. As research continues, images generated by GANs become more and more realistic and clear, and training procedure of GANs become more and more stable [9], [10]. Besides being an image generation technique, GANs can also be used for density estimation. The generator

part of a GAN is an implicit probability distribution model, and it will converge to a good estimator of the data distribution after training. In this work, we solve the inference problem with the probability density estimated by a GAN.

Figure 1 provides an illustration of how our image restoration method works. There are four dashed boxes from left to right in Figure 1, corresponding to four different phases of an image capture and restoration process. Images in the first dashed box are original images, which are clear and undegraded. These images undergo a series of degradation in the second dashed box, and then are captured by our camera. In the image restoration process, we hope to estimate the original images with the degraded images we captured. As we pointed out before, every degraded image could be degraded from many possible original images. To be more precise, there is a particular subset of the original image manifold for any degraded image, and all image samples on the submanifold could be degraded to the given degraded image. Images in the third dashed box are those samples on the submanifold, and they are arranged in ascending order of log-likelihood from left to right. Images marked by yellow boxes are samples with the highest log-likelihood in their group, and they are placed in the last dashed box as restoration outputs. Overall, the contributions of this work are mainly in two aspects:

- 1) We propose a general method for various image restoration problems. In the method, we explicitly use density information estimated by a GAN, an implicit model; and we directly solve the image restoration problem, an inference problem, with a GAN, a generative model. To the best of our knowledge, our work is the first to do those two things.
- 2) We propose a new algorithm to solve the optimization problem in our method. It is a first-order iterative algorithm for constrained problems, and it works well even for problems with highly nonlinear objectives and constraints. These features make it especially suited to neural network related constrained optimization problems.

II. RELATED WORK

The most similar works to ours are found in [11], [12]. [11] propose an image inpainting method, which can generate missing content with a trained GAN. They search in the latent




Fig. 1: An illustration of our image restoration method.

space of the GAN for the image which is close to the corrupted image, and use the discriminator loss as an indicator of how realistic their restoration is. [12] then improve their theory and apply it to various image restoration problems. Unfortunately, there is a major theoretical flaw in their method. [6] prove that the discriminator is unable to identify how realistic an input is after several steps of training, if the GAN has enough capacity. During the training, the information of the data distribution gradually transfer from the discriminator to the generator. Ideally, the generator will have all the information of the data distribution while the discriminator will have none. Even worse, [12] ignore a term $|\frac{\partial z}{\partial x}|$ intentionally in their Eq. (5), because they think it is intractable. We will demonstrate in the next section that $|\frac{\partial z}{\partial x}|$ directly determines the density of data space. In other words, [12] pays attention to the trivial and ignores the vital. That is why we use the generator instead of discriminator to measure how realistic the restoration is. Another difference between their work and ours is that we use a more radical strategy of optimization. They simply add their image prior term to their distortion. This will lead to a compromise between plausibility and visual quality of restoration. However, we choose the most probable image only from images which could degrade to the input. This makes our restorations more plausible while still keeping them similar to reality.

The maximum a posteriori (MAP) has existed for a long time as a classic estimation method [13], [14]. But before GANs, people do not have a probability density model which is good enough to describe the distribution of images. After GANs attained great success in image generation, researchers started to use them in image restoration tasks to get more realistic results [15], [16]. [17] and [18] try to use the MAP estimation on GANs to solve the image super-resolution problem. However, they only use the MAP estimation implicitly and indirectly, while our method uses it explicitly and directly. We suspect that all methods that do implicit MAP estimation on GANs would require redesigning or retraining when the image restoration task changes, and this makes implicit methods not as general as our explicit method.

[19] is another work which is seemingly similar to ours, but

they are actually quite different. They use a randomly-initialized neural network as a prior to solve image restoration problems. The prior in their method is elaborate, neural network related but still handcrafted, while in our method the prior is learned from data. So our data-driven prior has better adaptability to specific image distribution.

III. MAXIMUM A POSTERIORI ON A SUBMANIFOLD

A. Formulation

Consider a general image degradation model $\tilde{x} = F(x, \Omega)$, where x , \tilde{x} , and Ω represent the original image, the degraded image, and the parameters of the degradation model, respectively. The image degradation function F is a deterministic function. That means, given an original image x and a particular set of parameters Ω , the image degradation model will always produce the same degraded image \tilde{x} .

Our goal is to get a reasonable estimate of x with given \tilde{x} and F . In this paper, we use the maximum a posteriori probability (MAP) estimate of x as the restoration of \tilde{x} . Compared to MSE-based method, MAP estimate of x is perceptually more convincing. We can perform inference by maximizing the posterior $p(x, \Omega | \tilde{x})$:

$$\begin{aligned} \{\hat{x}, \hat{\Omega}\} &= \operatorname{argmax}_{x, \Omega} p(x, \Omega | \tilde{x}) \\ &= \operatorname{argmax}_{x, \Omega} \frac{p(\tilde{x} | x, \Omega) p(x | \Omega) p(\Omega)}{p(\tilde{x})} \end{aligned} \quad (1)$$

where \hat{x} and $\hat{\Omega}$ represent MAP estimate of x and Ω . Note that $p(\tilde{x})$ is always positive and does not depend on x and Ω , and typically we assume that x and Ω are independent. Therefore,

$$\{\hat{x}, \hat{\Omega}\} = \operatorname{argmax}_{x, \Omega} p(\tilde{x} | x, \Omega) p(x) p(\Omega) \quad (2)$$

Note that $\tilde{x} = F(x, \Omega)$ is a deterministic function, i.e., $p(\tilde{x} | x, \Omega) = \delta(\tilde{x} - F(x, \Omega))$. Therefore, the estimation is equivalent to

$$\begin{aligned} \{\hat{x}, \hat{\Omega}\} &= \operatorname{argmax}_{x, \Omega} p(x) p(\Omega) \\ \text{s.t. } &\|\tilde{x} - F(x, \Omega)\| = 0 \end{aligned} \quad (3)$$

Here we write $p(\mathbf{x})$ more specifically as $p_r(\mathbf{x})$, which stands for the probability density of real data distribution. We can estimate $p_r(\mathbf{x})$ with the generator part of a trained GAN, which is an implicit probability distribution model with distribution $p_G(\mathbf{x})$. The trained generator G represents a mapping from latent space of \mathbf{z} to data distribution of original image \mathbf{x} , i.e., $p_r(\mathbf{x}) = p_G(\mathbf{x})$, and $p_G(\mathbf{x})$ is a probability density function implicitly defined by $\mathbf{x} = G(\mathbf{z})$, where \mathbf{z} is typically sampled from some simple distribution, such as the uniform distribution or the normal distribution. Assuming $G : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is an injective function, the estimation is equivalent to

$$\{\hat{\mathbf{z}}, \hat{\boldsymbol{\Omega}}\} = \operatorname{argmax}_{\mathbf{z}, \boldsymbol{\Omega}} p_G(G(\mathbf{z}))p(\boldsymbol{\Omega}) \quad (4)$$

$$\text{s.t. } \|\tilde{\mathbf{x}} - F(G(\mathbf{z}), \boldsymbol{\Omega})\| = 0$$

$$\text{and } \hat{\mathbf{x}} = G(\hat{\mathbf{z}}) \quad (5)$$

Generally the dimension of vector space of \mathbf{z} is far lower than the dimension of vector space of \mathbf{x} . Note that $p_G(\mathbf{x})$ is nonnegative if and only if \mathbf{x} is on the low dimensional manifold \mathcal{M} defined by $\mathbf{x} = G(\mathbf{z})$, we can replace the probability density on the original space $p_G(G(\mathbf{z}))$ in Eq. (4) by the probability density on the manifold $p_{\mathcal{M}}(\mathbf{z})$, and end up with the same estimation result $\hat{\mathbf{z}}$. According to [20], the probability density on the manifold can be calculated by

$$p_{\mathcal{M}}(\mathbf{z}) = \frac{p(\mathbf{z})}{\sqrt{\det \operatorname{Gram}\left(\frac{\partial G}{\partial z_1}, \dots, \frac{\partial G}{\partial z_n}\right)}} \quad (6)$$

where Gram represents the Gram matrix, and $\sqrt{\det \operatorname{Gram}\left(\frac{\partial G}{\partial z_1}, \dots, \frac{\partial G}{\partial z_n}\right)}$ is the volume of the parallelotope spanned by the vectors $\left(\frac{\partial G}{\partial z_1}, \dots, \frac{\partial G}{\partial z_n}\right)$, so the square root of the Gram determinant can serve as a local scale factor. It has an effect similar to the Jacobian determinant, but we can only use the Gram determinant here because G is a function from \mathbf{R}^n to \mathbf{R}^m , and generally n is much less than m .

The Gram matrix can be simply calculated by $\operatorname{Gram}\left(\frac{\partial G}{\partial z_1}, \dots, \frac{\partial G}{\partial z_n}\right) = \mathbf{V}^T \mathbf{V}$, where \mathbf{V} is an $m \times n$ matrix, whose entries are given by $V_{ij} = \frac{\partial x_i}{\partial z_j}$. Therefore, Eq. (4) is equivalent to

$$\{\hat{\mathbf{z}}, \hat{\boldsymbol{\Omega}}\} = \operatorname{argmax}_{\mathbf{z}, \boldsymbol{\Omega}} \frac{p(\mathbf{z})p(\boldsymbol{\Omega})}{\sqrt{\det \mathbf{V}^T \mathbf{V}}} \quad (7)$$

$$\text{s.t. } \|\tilde{\mathbf{x}} - F(G(\mathbf{z}), \boldsymbol{\Omega})\| = 0$$

To solve the estimation problem efficiently, we represent probabilities in Eq. (7) in logarithmic space, i.e.,

$$\log \frac{p(\mathbf{z})p(\boldsymbol{\Omega})}{\sqrt{\det \mathbf{V}^T \mathbf{V}}} = -\frac{1}{2} \log \det \mathbf{V}^T \mathbf{V} + \log p(\mathbf{z}) + \log p(\boldsymbol{\Omega}) \quad (8)$$

Matrix $\mathbf{V}^T \mathbf{V}$ is a positive-definite matrix, so we can use Cholesky decomposition to calculate $\log \det \mathbf{V}^T \mathbf{V}$ efficiently, i.e.,

$$\log \det \mathbf{V}^T \mathbf{V} = 2 \operatorname{tr}(\log(\operatorname{chol}(\mathbf{V}^T \mathbf{V}))) \quad (9)$$

Finally we deduce a set of expressions which can be calculated directly, and their final outcome $\hat{\mathbf{x}}$ is the restored image we want, i.e.,

$$\{\hat{\mathbf{z}}, \hat{\boldsymbol{\Omega}}\} = \operatorname{argmax}_{\mathbf{z}, \boldsymbol{\Omega}} -\operatorname{tr}(\log(\operatorname{chol}(\mathbf{V}^T \mathbf{V}))) + \log p(\mathbf{z}) + \log p(\boldsymbol{\Omega}) \quad (10)$$

$$\text{s.t. } \|\tilde{\mathbf{x}} - F(G(\mathbf{z}), \boldsymbol{\Omega})\| = 0$$

$$\text{and } \hat{\mathbf{x}} = G(\hat{\mathbf{z}}) \quad (11)$$

Note that $(G(\mathbf{z}), \boldsymbol{\Omega})$ form a low dimensional manifold which is embedded in the space of $(\mathbf{x}, \boldsymbol{\Omega})$, and the feasible solutions of Eq. (10) is on a subset of the manifold, which is defined by $\|\tilde{\mathbf{x}} - F(G(\mathbf{z}), \boldsymbol{\Omega})\| = 0$. So our method basically makes a MAP estimate on a submanifold.

Figure 2 is a toy example to show the basic idea of our formulation in a very visible way. Suppose there is a grayscale original image \mathbf{x} , which has only three pixels. Then it is downsampled to only one pixel during the image capture process, and our task is to estimate \mathbf{x} with the one pixel image we captured. Suppose we have trained a GAN as an implicit model of data distribution of \mathbf{x} . More specifically, the generator of the trained GAN represents a mapping from its input noise \mathbf{z} to data distribution of \mathbf{x} . The left part of Figure 2 describes the two dimensional latent space of \mathbf{z} . We use the saturation of orange color to represent probability density level, i.e., a thicker orange color means a higher probability density. So the uniform orange color in the latent space means that the input noise \mathbf{z} is sampled from a uniform distribution.

Then the two dimensional vector \mathbf{z} is mapped to three dimensional space of image \mathbf{x} by the generator of the trained GAN, and the big orange square in the latent space of \mathbf{z} is transformed into a twisted torus in the three dimensional data space of \mathbf{x} , which is described in the right part of Figure 2. Some areas in space of \mathbf{z} expand during the transformation, while other areas shrink. We can find this out by comparing the red and blue quadrilateral between the latent and data space. Therefore, the probability density on the torus is no longer uniform. The orange colors of the expanded areas become lighter, and the colors of the shrunken areas become thicker. Quantitatively speaking, the square root of the Gram determinant in Eq. (6) is the local area scale factor of the mapping, and its inverse, of course, is the local density scale factor.

The pale yellow plane in the data space represents the constraint in the toy example. All points on the plane would exactly be downsampled to the one pixel image we captured. So the intersection curve of the plane and the torus is the submanifold we are looking for, and that white curve is the feasible set of the toy problem. In this problem, $p(\mathbf{z})$ is a constant in the domain, and degradation parameters $\boldsymbol{\Omega}$ does not exist at all. According to Eq. (7), what we need to do is to maximize the inverse of the square root of the Gram determinant on the submanifold. In other words, the point with the thickest orange color on the intersection curve is the restored image $\hat{\mathbf{x}}$, the MAP estimate on the submanifold. We

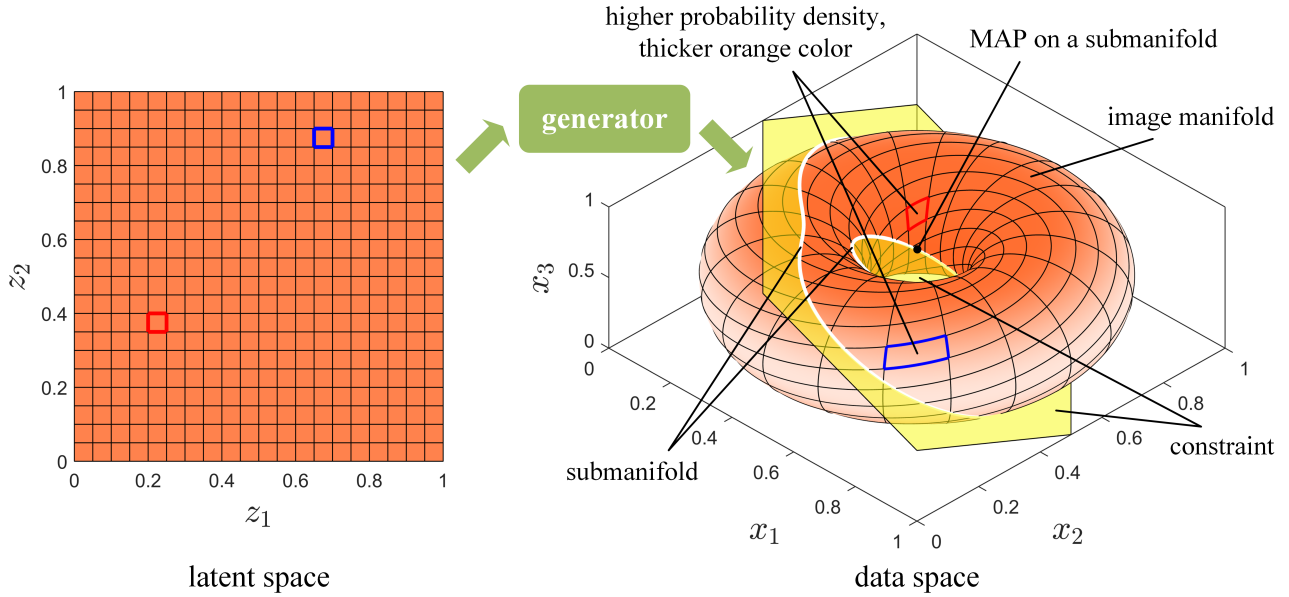


Fig. 2: A toy example to show the basic idea of our formulation.

can find out that the method is both intuitive and rational for this toy example.

B. Optimization Algorithm

We propose a new optimization algorithm to solve Eq. (10). Note that the objective function and the equality constraint in Eq. (10) are both highly nonlinear, so gradient-based method seems a natural choice for the problem. Our algorithm is inspired by Projected Gradient Descent Method.

To solve an unconstrained problem with ordinary Gradient Descent Method, we take small steps in the direction of the negative gradient. To solve a constrained problem, we can try to use Projected Gradient Descent Method, take a small step as usual and then project variables back onto the feasible set. But unfortunately, Projected Gradient Descent Method is only valid for problems with a very simple feasible set, such as a solution set of linear equations, a simple polyhedron or a simple cone. If constraints of a problem are too complex, like the constraint in Eq. (10), it is very hard to project variables back onto the feasible set.

To overcome this shortage, we propose a new optimization algorithm called Quasi Projected Gradient Descent Method. In our algorithm, the gradient information is not only used to improve the objective function, but helps to satisfy the constraints as well. Consider the standard form of a continuous optimization problem,

$$\begin{aligned} & \underset{\mathbf{u}}{\text{minimize}} && f(\mathbf{u}) \\ & \text{s.t.} && h_i(\mathbf{u}) = 0, \quad i = 1, \dots, m \\ & && h_j(\mathbf{u}) \leq 0, \quad j = m+1, \dots, m+p \end{aligned} \quad (12)$$

where $f, h_i, h_j : \mathbf{R}^n \rightarrow \mathbf{R}$, and they are all highly nonlinear. Algorithm 1 is the proposed algorithm for the problem.

To solve Eq. (10) with Algorithm 1, we only need to set $\mathbf{u} = \{\hat{\mathbf{z}}, \hat{\Omega}\}$, objective function $f(\mathbf{u}) =$

Algorithm 1 Quasi Projected Gradient Descent Method

Input: objective function $f(\mathbf{u})$, equality constraints $h_i(\mathbf{u})$ and inequality constraints $h_j(\mathbf{u})$

Parameters: step size η_{\parallel} and η_{\perp} , positive factors c_i and c_j , number of iterations n , small positive constant ϵ for numerical stability, initial guess \mathbf{u}_0

Output: local optimum \mathbf{u}_n

Let $h(\mathbf{u}) = \sum_{i=1}^m c_i \cdot \|h_i(\mathbf{u})\|^2 + \sum_{j=m+1}^{m+p} c_j \cdot H(h_j(\mathbf{u})) \cdot \|h_j(\mathbf{u})\|^2$, where H represents the Heaviside step function

for $i = 1$ **to** n **do**

$$\mathbf{g}_f = \nabla f(\mathbf{u}_{i-1})$$

$$\mathbf{g}_h = \nabla h(\mathbf{u}_{i-1})$$

$$\mathbf{g}_{\parallel} = \mathbf{g}_f - \frac{\mathbf{g}_f \cdot \mathbf{g}_h}{\mathbf{g}_h \cdot \mathbf{g}_h + \epsilon} \cdot \mathbf{g}_h$$

$$\mathbf{g}_{\perp} = \mathbf{g}_h$$

$$\mathbf{u}_i = \mathbf{u}_{i-1} - \eta_{\parallel} \cdot \mathbf{g}_{\parallel} \quad (\text{or use more advanced optimizer})$$

$$\mathbf{u}_i = \mathbf{u}_i - \eta_{\perp} \cdot \mathbf{g}_{\perp} \quad (\text{or use more advanced optimizer})$$

end for

return \mathbf{u}_n

$-(-\text{tr}(\log(\text{chol}(\mathbf{V}^T \mathbf{V}))) + \log p(\mathbf{z}) + \log p(\Omega))$, and the only equality constraint function $h_1(\mathbf{u}) = \|\hat{\mathbf{x}} - F(G(\mathbf{z}), \Omega)\|$.

In the proposed algorithm, we first define an overall constraint function $h(\mathbf{u}) : \mathbf{R}^n \rightarrow \mathbf{R}_{\geq 0}$, and the feasible set of the optimization problem is the region where $h(\mathbf{u}) = 0$. In each iteration of the algorithm, we calculate the gradients of $f(\mathbf{u})$ and $h(\mathbf{u})$ at \mathbf{u}_{i-1} . If we take a small step in the direction of the negative \mathbf{g}_f , the value of $f(\mathbf{u})$ will decrease a little bit, but it may have an unwanted impact on the value of $h(\mathbf{u})$. In order to avoid this problem, we calculate \mathbf{g}_{\parallel} , the tangential component of \mathbf{g}_f on the isocontour of $h(\mathbf{u}_{i-1})$, which can be calculated by vector rejection of \mathbf{g}_f on \mathbf{g}_h . In each iteration, we actually take a small step in the direction of the negative

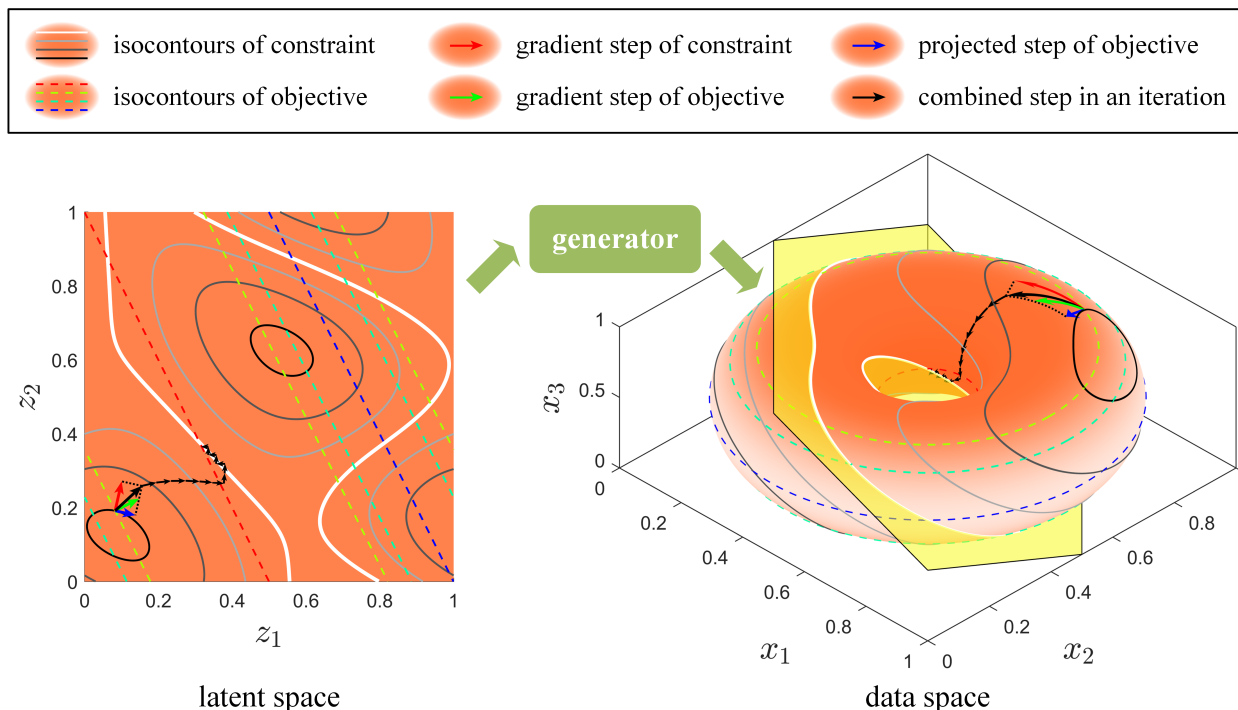


Fig. 3: A toy example to show how our Quasi Projected Gradient Descent Method works.

\mathbf{g}_{\parallel} , the value of $f(\mathbf{u})$ will still decrease, while it has almost no impact on the value of $h(\mathbf{u})$. We also take a small step in the direction of the negative \mathbf{g}_{\perp} , i.e., \mathbf{g}_h itself, which is perpendicular to the isocontour of $h(\mathbf{u}_{i-1})$. Repeat these steps, and the sequence \mathbf{u} will hopefully converge to the desired optimal solution. Behaviors of our Quasi Projected Gradient Descent Method is similar to behaviors of the original Projected Gradient Descent Method. Consider a point \mathbf{u} which is very close to the feasible region. The summation of two moves against \mathbf{g}_{\parallel} and \mathbf{g}_{\perp} is actually an inaccurate Projected Gradient Descent. That is why we name our method as Quasi Projected Gradient Descent Method.

Here we use the same toy example we used in Section III-A, to show how our Quasi Projected Gradient Descent Method works. In Figure 3, solid curves in black and white are isocontour of constraint function h . The whiter the curve, the lower value of h it corresponds; Dashed lines in color are isocontour of objective function f . The redder the line, the lower value of f it corresponds. Note that the white solid curve is the feasible set of the toy problem, so intersection points of the white solid curve and the red dashed line in the latent space is $\hat{\mathbf{z}}$ in Eq. (10), while the intersection points in the data space is $\hat{\mathbf{x}}$ in Eq. (11).

Our iterative optimization algorithm starts from the bottom left corner of the latent space. The red vector is a gradient step of h . It is pointing towards the direction of the negative \mathbf{g}_h , and is perpendicular to the black solid curve, an isocontour of h . The green vector is a gradient step of f . It is pointing towards the direction of the negative \mathbf{g}_f , and is perpendicular to the yellow dashed line, an isocontour of f . The blue vector

is a projected gradient step. It is pointing towards the direction of the negative \mathbf{g}_{\parallel} , and is the tangential component of the green vector on the black solid curve, which can be calculated by vector rejection of the green vector on the red vector. We only plot green, red and blue vector for the first iteration to keep Figure 3 clean and easy to understand. Black vectors are combined gradient steps, which are vector sums of red and blue vectors. We move along these black vectors and we can find out that our optimization algorithm reaches a desired solution quickly.

IV. EXPERIMENTS

In this Section, we use MNIST dataset [21] to test our image restoration method. The dataset is divided in 50k for the training set, 10k for each of the validation and test set. We use a WGAN-GP [22] trained on the training set as the density estimation model. The architecture of the WGAN-GP we used is shown in Table III and Table IV, and we add a L2 weight decay term with decay parameter of 0.001 on the generator loss to prevent over-fitting. The network we used is very simple, but it is enough to prove the effectiveness of our method.

We use four different kinds of degradation to test the generality of our method. The first three kinds of degradation are relatively simple. They are 7× downsampling, making a 14×14 square hole in the center of the image, and adding Gaussian white noise with a standard deviation of 1.0, respectively. The last kind of degradation is a composition of a series of degradation in order, which are (a) adding linear motion blur by at most 14 pixels in any direction, (b) 4× downsampling,

TABLE I: Quantitative comparison with other general image restoration methods using PSNR(dB) and SSIM.

	Downsample		Hole		Noise		Composition	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Total Variation	12.88	0.4556	11.06	0.5928	11.94	0.1736	12.64	0.2559
[12]	13.23	0.6612	12.13	0.6205	11.99	0.5910	12.79	0.6434
Ours	17.02	0.8287	14.63	0.7815	14.47	0.7238	14.73	0.7403

TABLE II: Visual comparison with other general image restoration methods.

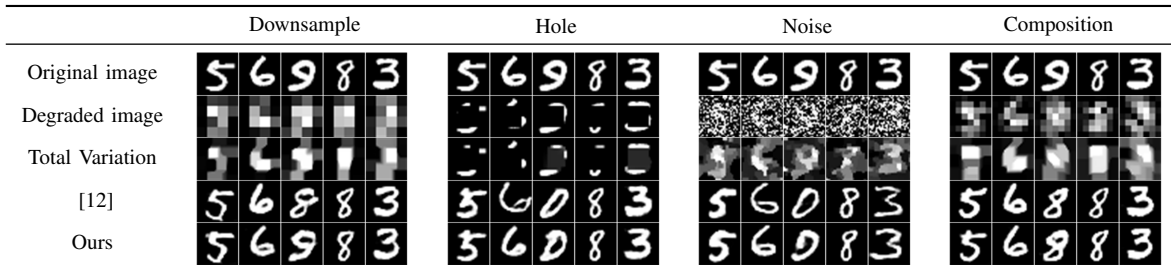


TABLE III: Architecture of the generator

	Kernel size	Output shape
z	-	16
Linear, tanh	-	$64 \times 4 \times 4$
Deconv, tanh	5×5	$32 \times 7 \times 7$
Deconv, tanh	5×5	$16 \times 14 \times 14$
Deconv, sigmoid	5×5	$1 \times 28 \times 28$

TABLE IV: Architecture of the discriminator

	Kernel size	Output shape
$G(z)$	-	$1 \times 28 \times 28$
Conv, LeakyReLU	5×5	$16 \times 14 \times 14$
Conv, LeakyReLU	5×5	$32 \times 7 \times 7$
Conv, LeakyReLU	5×5	$64 \times 4 \times 4$
Linear	-	1

(c) adding uniform noise between -0.05 and 0.05, (d) randomly removing 10% of the pixels.

We use two independent ADAM optimizers [23] with g_{\parallel} and g_{\perp} respectively in the Quasi Projected Gradient Descent Method. For all four kinds of degradation, we run the algorithm with the same settings. Settings for both ADAM optimizer are learning rate $\alpha = 0.01$ (decayed linearly to 0), $\beta_1 = 0.9$, $\beta_2 = 0.99$, and number of iterations $n = 500$.

In the experiments, we compare our method with two other general image restoration methods. The first is Total Variation (TV) [24], a traditional method; and the second is [12], a GAN based state-of-the-art approach. We use the SSIM index [25] and PSNR value as quantitative metrics for these restoration methods, and the results are shown in Table I. We also present some restoration images in Table II without cherry-picking. We can find out that our general image restoration method is better than both baseline methods by large margins. This is due to the more accurate prior information of images and the more radical strategy of optimization in our method.

V. CONCLUSIONS AND FUTURE WORK

We propose a general image restoration method in this work. Compared with traditional image restoration algorithms, our method is much more powerful. Image restoration is an inherently ill-posed problem, so additional prior knowledge is needed. In our method, we use all prior knowledge of original images, i.e., the probability distribution of degradation, i.e., the degradation model itself. Traditional image restoration like Total Variation, by contrast, just uses a small part of the prior, typically some statistical properties. Besides, unlike our method, there is usually no guarantee that an output restoration from a traditional method can be degraded back accurately to its input. This makes restorations from a traditional method less plausible than restorations from our method.

For future work, we think our method can be straightforwardly extended to other domains which GANs are gifted in, such as video, audio and language. We will try to solve restoration problems and other inference problems in these domains with our paradigm. The convergence and other properties of the Quasi Projected Gradient Descent Method would be interesting as well.

REFERENCES

- [1] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE signal processing magazine*, vol. 20, no. 3, pp. 21–36, 2003.
- [2] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Transactions on image processing*, vol. 17, no. 1, pp. 53–69, 2008.
- [3] C. Guillemot and O. Le Meur, "Image inpainting: Overview and recent advances," *IEEE signal processing magazine*, vol. 31, no. 1, pp. 127–144, 2014.
- [4] J. Pan, D. Sun, H. Pfister, and M.-H. Yang, "Blind image deblurring using dark channel prior," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 1628–1636.
- [5] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.

- [6] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [7] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [8] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," *arXiv preprint arXiv:1809.11096*, 2018.
- [9] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," in *Advances in Neural Information Processing Systems*, 2016, pp. 2234–2242.
- [10] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," *arXiv preprint arXiv:1701.07875*, 2017.
- [11] R. A. Yeh, C. Chen, T.-Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do, "Semantic image inpainting with deep generative models," in *CVPR*, vol. 2, no. 3, 2017, p. 4.
- [12] R. A. Yeh, T. Y. Lim, C. Chen, A. G. Schwing, M. Hasegawa-Johnson, and M. Do, "Image restoration with deep generative models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6772–6776.
- [13] M. E. Tipping and C. M. Bishop, "Bayesian image super-resolution," in *Advances in neural information processing systems*, 2003, pp. 1303–1310.
- [14] P. Campisi and K. Egiazarian, *Blind image deconvolution: theory and applications*. CRC press, 2016.
- [15] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," *arXiv preprint*, 2017.
- [16] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, no. 2, 2017, p. 7.
- [17] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *CVPR*, vol. 2, no. 3, 2017, p. 4.
- [18] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, "Amortised map inference for image super-resolution," *arXiv preprint arXiv:1610.04490*, 2016.
- [19] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," *arXiv preprint arXiv:1711.10925*, 2017.
- [20] X. Pennec, "Probabilities and statistics on riemannian manifolds: A geometric approach," Ph.D. dissertation, INRIA, 2004.
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *Advances in Neural Information Processing Systems*, 2017, pp. 5767–5777.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, "An augmented lagrangian approach to the constrained optimization formulation of imaging inverse problems," *IEEE Transactions on Image Processing*, vol. 20, no. 3, pp. 681–695, 2011.
- [25] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli *et al.*, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.